

CS290F - Paper Reviews for 2010.01.27

In Depth Review for: Banerjee, Bhattacharjee, "A Comparative Study of Application Layer Multicast Protocols", unpublished.

When presented with a paper that has never been accepted for publication, a reviewer should be ready for a paper that may fall anywhere in a very wide spectrum of both technical and compositional quality; in the case of the unpublished paper by Banerjee and Bhattacharjee, the paper falls somewhere above poor but below average on both scales: from a technical perspective, the highlights of the purpose and implementation of the protocols are adequately conveyed, but fine details and comparative evaluation are conspicuous in their absence; From a compositional perspective, the paper attempts to follow the standard comparison format, progressing through an introduction, discussion, and analysis, but poor grammar, awkward transitions, useless filler text, and typos are abundant. The remainder of this review will attempt to remain focused on the technical merits and flaws in the paper, but in many cases the composition and writing style of the paper significantly impact the technical discussion presented therein, and therefore merit at least some moderate mention.

The abstract of the paper begins with a fairly solid summary of the basic concept behind application layer multicast, but fails to provide any indication of how that idea is going to be dealt with in the contents of the paper, beyond a weak "...comparison of performance and application...". The abstract of a paper should explain at a high level how various ideas are going to be compared, contrasted, and evaluated; this abstract leaves the reader with no more information than is provided by the title itself, leaving little reason for a conference reviewer to continue reading the paper.

A recurring theme throughout the paper is the unsupported claim. An example of one such claim presents itself early in the introduction of the paper: "IP Multicast is the most efficient way to perform group data distribution, as it is able to reduce packet replication on the wide-area network to the minimum necessary." Now, this particular statement is less egregious than other unsupported claims in the paper since at least a supporting claim is attached to it. However, the claim itself would carry significantly more weight with even a simple citation attached to it from nearly any paper that addresses nearly any implementation of IP multicast; as is, the lack of a supporting citation or argument justifying the supporting claim leaves the reader wondering exactly the opposite: if there is some other, more efficient, packet delivery system (e.g. consider a system that somehow avoids much of the IP multicast group registration and maintenance).

As if to drive home the point, the very next paragraph includes a similarly unsupported claim: "[U]nlike unicast addresses, these multicast group addresses are not easily aggregatable [sic]." Even more than the last example, this quote begs for a supporting citation -- after all, the paper's authors must have come across this idea somewhere, so why not add the citation here so the reader can verify the claim directly?. The remainder of the paper continues to make unsupported claims, particularly when discussing the *qualities* of an algorithm or concept -- in fact, nearly all of the citations in the paper are

limited to direct references to a protocol implementation. Therefore, to minimize the beating of a dead horse, the remainder of this review will avoid pointing out unsupported claims in all but the most severe of omissions.

The paper first delves into some technical content near the middle of the introduction section as it begins to discuss the metrics by which the paper will evaluate the protocols that are compared by the study. Unfortunately, this is also where the paper begins to make technical and compositional errors that are somewhat more critical than mere omissions of citations. In particular, the paper presents merely two metrics by which a protocol is to be evaluated: the "quality of the data path" (which itself is divided into two further metrics: "stress" and "stretch") and the "control overheads". Now, it may be the case that these particular metrics are the "best" metrics for evaluating application layer multicast protocols, but the paper makes no substantial defense for these metrics, and only recognizes other metrics in a hand-waving manner. Several other metrics that seem missing include "memory consumption", "physical path length" (which differs from stretch in that stretch measures length on the application layer overlay network, while physical path length measures the number of actual hops between routers), and/or "average delay" (in terms of nodes and path length). This small sample group of metrics seem particularly interesting when comparing application layer protocols, but are neither mentioned nor discussed anywhere in the paper.¹ (The authors of the comparison study are free to select the metrics they deem "best" for the study, but they have an obligation to defend that selection.)

Figures (including tables and illustrations) in a paper should help a paper further demonstrate a point or help present data in a meaningful way. Additionally, figures ought to be self-explanatory. That is, a reader should be able to look at a figure and its caption and labels and understand the meaning and purpose of the figure. To pick a particularly poor example in this paper, looking solely at figure 1, a reader cannot tell much, if anything, about which application layer multicast mechanisms are in use, until they have derived the algorithm by following the message-flow arrows provided. By simply providing more useful labels, the figure would instantly become more clear.

From the opposite point of view, a figure should not be the *only* explanation of a concept in a paper. In this case, the end of the introduction to the paper contains prose that (briefly) walks the reader through figure 1, but the concepts themselves are otherwise left for the reader to derive (e.g., the discussion of panel 3 is a single sentence that states the overlay is a configuration somewhere between the other two configurations). Had the authors taken the time to discuss the different concepts involved in more detail, and then used the figure to further explain the concept, the reader would be left with a complete understanding of the metrics and overlay concepts referred to throughout the remainder of the paper.

Following the introductory section (which is in reality a combination of what are more commonly separated into separate introduction and background sections), the paper presents three separate approaches to application layer multicast: a "mesh-first approach", a "tree-first approach", and an "implicit approach". (After these approaches, an "other protocols" section is included that briefly mentions a handful of protocols that the authors didn't see fit to include in one of their three sections for one reason or another.) Each of these sections provide a reasonable explanation of the basic concept of the approach, and

1. The critical reader might accuse this review of the same hand-waving demonstrated by the paper in question. However, the standard to which this review is held (in terms of this list) should be somewhat weaker, since, unlike the paper in question, this review is not intended to be a comparative study of application layer multicast protocols, and isn't intended to be submitted as a representative comparison to a conference.

one to three implemented protocols are discussed in relation to each approach. These sections are fairly unremarkable in that they a) present the ideas in a succinct manner, but b) don't delve terribly deep into any of the particular ideas.

Finally, in the seventh section of the paper, we arrive at the meat of the paper: the comparative study itself... only to find that the paper really provides little to no comparison of the different approaches, other than a small, sparsely populated table of values that describe the metrics presented earlier in terms of the number of nodes in the multicast layer. A couple of generally unsupported claims are made, attempting to describe where each of the approaches will fit best, but no meaningful analysis of any of the approaches is presented, and no supporting citations or arguments are given for any of the claims made. The entire "comparative study" section is less than a single page of the just over eight pages of text in the paper. Unsurprisingly, the conclusion of the paper is merely some filler text -- after all, if there was no actual comparison performed, at what conclusions could the paper really arrive?